

CLAIMS

We claim:

1. A method to transfer data from a sending application to a receiving application in a computer environment having a plurality of data transfer modes and at least one system area network, the method comprising:

detecting a transfer behavior of the receiving application; and

switching the data transfer mode to one of the plurality of data transfer modes based upon the transfer behavior.

2. The method of claim 1 further comprising the step of transferring data in one of the data transfer modes.

3. The method of claim 1 wherein the step of detecting a transfer behavior of the receiving application comprises the steps of:

transferring in an initial transfer mode from the sending application to the receiving application a pre-selected number of initial data blocks each having a size greater than a pre-selected threshold size; the transferring of each of the initial data blocks in the initial transfer mode including:

(i) receiving a transfer request for said each initial data block sent by the sending application;

(ii) sending a pilot message for informing the receiving application that said each initial data block is available to be sent;

1. A method to transfer data from a sending application to a receiving application in a computer environment having a plurality of data transfer modes and at least one system area network, the method comprising:
 detecting a transfer behavior of the receiving application; and
 switching the data transfer mode to one of the plurality of data transfer modes based upon the transfer behavior.
 2. The method of claim 1 further comprising the step of transferring data in one of the data transfer modes.
 3. The method of claim 1 wherein the step of detecting a transfer behavior of the receiving application comprises the steps of:
 transferring in an initial transfer mode from the sending application to the receiving application a pre-selected number of initial data blocks each having a size greater than a pre-selected threshold size; the transferring of each of the initial data blocks in the initial transfer mode including:
 (i) receiving a transfer request for said each initial data block sent by the sending application;
 (ii) sending a pilot message for informing the receiving application that said each initial data block is available to be sent;

(iii) detecting whether the receiving application posts a receive buffer when receiving the transfer request for said each initial data block or when the receiving application is informed that said each data block is available to be sent; and

(iv) determining whether the receive buffer exceeds the threshold size.

4. The method of claim 3 further comprising the steps of:

if the receiving application posts a receive buffer exceeding the threshold size:

(i) locating the receive buffer posted by the receiving application;

(ii) transferring the initial data block to the receiving application by direct memory access to the receive buffer; and

wherein if for each of the initial data blocks the receiving application posts a receive buffer exceeding the threshold size when sending a transfer request for said each initial data block, switching the data transfer mode for subsequent data blocks having sizes greater than the threshold size to a large-receive transfer mode.

5. The method of claim 4 further comprising the step of transferring subsequent data blocks having sizes greater than the threshold size in a large-receive transfer mode.

6. The method of claim 5 wherein the step of transferring subsequent data blocks having sizes greater than the threshold size in a large-receive transfer mode comprises the step of:

for each subsequent data block:

(i) locating a receive buffer posted by the receiving application;

(ii) transferring said each subsequent data block to the receiving application by direct memory access to the receive buffer.

7. The method of claim 1 wherein the step of switching the data transfer mode based upon the transfer behavior comprises the step of:

if the transfer behavior of the receiving application is to post a receive buffer exceeding a threshold size when the receiving application is informed that a data block is available to be sent: switching the data transfer mode to a small-receive-large-receive transfer mode.

8. The method of claim 7 further comprising the step of:

if the data transfer mode is a small-receive-large-receive transfer mode:

transferring subsequent data blocks having sizes greater than the threshold size in a small-receive-large-receive transfer mode.

9. The method of claim 1 wherein the step of detecting the transfer behavior comprises the step of detecting if the receiving application posts receive buffers smaller than a threshold size when the receiving application receives a transfer request or when the receiving application is informed that said each data block is available to be sent.

10. The method of claim 9 wherein the step of switching the data transfer mode comprises the step of switching the data transfer mode to a small-receive transfer mode if the transfer behavior is that the receiving application posting receive buffers smaller than the threshold size when the receiving application receives the transfer request or when the receiving application is informed that said each data block is available to be sent

11. The method of claim 10 further comprising the step of sending the data in a plurality of messages, each message having a portion of the data for each subsequent data block, if the transfer mode is the small-receive transfer mode.

12. The method of claim 9 further comprising the steps of:

sending the data in at least one message having an initial buffer size;

if a number of fragmentations of data exceeds a threshold level during a time when a predetermined number of messages is sent:

(a) sending a resize request message having a requested buffer size to the receiving application;

(b) receiving a new buffer size from the receiving application;

(c) sending data to the receiving application in a data size equal to the new buffer size.

13. The method of claim 12 further comprising the steps:

determining a largest size of data block being sent; and

setting the requested buffer size to the largest size.

14. The method of claim 12 further comprising the step of sending data to the receiving application in a data size equal to the initial buffer size if the receiving application posts receiving buffers having a size equivalent to the initial buffer size.

15. The method of claim 12 wherein steps (a) - (c) are performed when the number of fragmentations of data exceeds a threshold level of approximately ten percent of the messages sent.

16. The method of claim 1 wherein the step of detecting the transfer behavior comprises the step of detecting whether an initial data block and a subsequent data block has been transferred within a predetermined time.

17. The method of claim 16 wherein the step of switching the data transfer mode comprises the step of switching the data transfer mode to small receive mode if the transfer behavior is that the initial data block and the subsequent data block has not been transferred within the predetermined time.

18. The method of claim 17 further comprising the step of transferring one of the initial data block and the subsequent data block in messages if the data transfer mode has been switched to small receive mode, each message containing a portion of one of the initial data block and the subsequent data block.

19. A computer-readable medium having computer-executable instructions for transferring data from a sending application to a receiving application in a computer environment having a plurality of data transfer modes the computer-executable instructions performing the steps of:

detecting a transfer behavior of the receiving application; and

switching the data transfer mode based upon the transfer behavior.

20. The computer-readable medium of claim 19 having further computer-executable instructions for performing the step comprising transferring data in one of the data transfer modes.

21. The computer-readable medium of claim 19 wherein the step of detecting a transfer behavior of the receiving application comprises the steps of:

transferring in an initial transfer mode from the sending application to the receiving application a pre-selected number of initial data blocks each having a size greater than a pre-selected threshold size; the transferring of each of the initial data blocks in the initial transfer mode including:

- (i) receiving a transfer request for said each initial data block sent by the sending application;
- (ii) sending a pilot message for informing the receiving application that said each initial data block is available to be sent;
- (iii) detecting whether the receiving application posts a receive buffer when receiving the transfer request for said each initial data block or when the receiving application is informed that said each data block is available to be sent; and
- (iv) determining whether the receive buffer exceeds the threshold size.

22. The computer-readable medium of claim 21 having further computer-executable instructions for performing the steps comprising:

if the receiving application posts a receive buffer exceeding the threshold size:

- (i) locating the receive buffer posted by the receiving application;
- (ii) transferring the initial data block to the receiving application by direct memory access to the receive buffer; and

wherein if for each of the initial data blocks the receiving application posts a receive buffer exceeding the threshold size when sending a transfer request for said each initial data

block, switching the data transfer mode for subsequent data blocks having sizes greater than the threshold size to a large-receive transfer mode.

23. The computer-readable medium of claim 22 having further computer-executable instructions for performing the steps comprising transferring subsequent data blocks having sizes greater than the threshold size in a large-receive transfer mode.

24. The computer-readable medium of claim 23 wherein the step of transferring subsequent data blocks having sizes greater than the threshold size in a large-receive transfer mode comprises the step of:

for each subsequent data block:

- (i) locating a receive buffer posted by the receiving application;
- (ii) transferring said each subsequent data block to the receiving application by direct memory access to the receive buffer.

25. The computer-readable medium of claim 19 wherein the step of switching the data transfer mode based upon the transfer behavior comprises the step of:

if the transfer behavior of the receiving application is to post a receive buffer exceeding a threshold size when the receiving application is informed that a data block is available to be sent: switching the data transfer mode to a small-receive-large-receive transfer mode.

26. The computer-readable medium of claim 25 having further computer-executable instructions for performing the step comprising

if the data transfer mode is a small-receive-large-receive transfer mode:

transferring subsequent data blocks having sizes greater than the threshold size in a small-receive-large-receive transfer mode

27. The computer-readable medium of claim 19 wherein the step of detecting the transfer behavior comprises the step of detecting if the receiving application posts receive buffers smaller than the threshold size when the receiving application receives the transfer request or when the receiving application is informed that said each data block is available to be sent.

28. The computer-readable medium of claim 27 having further computer-readable instructions for performing the steps comprising:

sending the data in at least one message having an initial buffer size;

if a number of fragmentations of data exceeds a threshold level during a time when a predetermined number of messages is sent:

(a) sending a resize request message having a requested buffer size to the receiving application;

(b) receiving a new buffer size from the receiving application;

(c) sending data to the receiving application in a data size equal to the new buffer size.

29. The computer-readable medium of claim 28 having further computer-executable instructions for performing the step comprising setting the requested buffer size to a largest size of data being sent.

30. The computer-readable medium of claim 28 having further computer-executable instructions for performing the step comprising sending data to the receiving application in a

data size equal to the initial buffer size if the receiving application posts receiving buffers having a size equivalent to the initial buffer size.

31. The computer-readable medium of claim 28 wherein steps (a) - (c) are performed when the number of fragmentations of data exceeds a threshold level of at least 30 during a time when approximately 300 messages are sent.

32. The computer-readable medium of claim 28 wherein steps (a) - (c) are performed when the number of fragmentations of data exceeds a threshold level of approximately ten percent of the messages sent.

33. The computer-readable medium of claim 27 wherein the step of switching the data transfer mode comprises the step of switching the data transfer mode to a small-receive transfer mode if the transfer behavior is that the receiving application posting receive buffers smaller than the threshold size when the receiving application receives the transfer request or when the receiving application is informed that said each data block is available to be sent

34. The computer-readable medium of claim 33 having further computer-executable instructions for performing the step comprising sending the data in a plurality of messages, each message having a portion of the data for each subsequent data block, if the transfer mode is the small-receive transfer mode.

35. The computer-readable medium of claim 19 wherein the step of detecting the transfer behavior comprises the step of detecting whether an initial data block and a subsequent data block has been transferred within a predetermined time.

36. The computer-readable medium of claim 35 wherein the step of switching the data transfer mode comprises the step of switching the data transfer mode to small receive mode if the transfer behavior is that the initial data block and the subsequent data block has not been transferred within the predetermined time.

37. The computer-readable medium of claim 36 having further computer-executable instructions for performing the step comprising transferring one of the initial data block and the subsequent data block in messages if the data transfer mode has been switched to small receive mode, each message containing a portion of one of the initial data block and the subsequent data block.

38. The computer-readable medium of claim 37 having further computer-executable instructions for performing the step comprising transferring one of the initial data block and the subsequent data block in messages if the data transfer mode has been switched to small receive mode, each message containing a portion of one of the initial data block and the subsequent data block.